Implication of Internet Traffic Characteristics for Network-Adaptive Distributed Systems

Jian Yuan and Kevin Mills
National Institute of Standards and Technology
NIST North 445
Gaithersburg, Maryland 20899 USA
Tel. 301-975-3618
Fax. 301-590-0932
jian.yuan@nist.gov and kmills@nist.gov

Abstract.

A growing number of distributed systems and network applications are deployed atop the Internet. Designers of such systems turn increasingly to adaptive mechanisms that measure network delays and alter processing according to conditions. In this paper, we argue that understanding not only the statistical properties of Internet traffic, but also the causes of those properties, can help designers of network-adaptive distributed systems to produce more effective measurement and adaptation strategies. We use simulation to illustrate effects on network traffic of some key factors: characteristics of offered traffic, end-to-end congestion-control mechanisms, link capacity, and network size. We apply wavelet-based analysis to illustrate the propagation of congestion over various timescales. In general, we find that network size has greater influence than other factors on network congestion. Our findings suggest that distributed systems should adapt to network conditions at a variable pace, rather than a constant pace.

Keywords.

network adaptation, network congestion, network latency, simulation, wavelet-based analysis

I. Introduction

To tune performance of distributed systems running over the global Internet, understanding Internet traffic characteristics should prove extremely useful. The Internet itself is essentially a big distributed system, where transport-layer traffic flows adapt themselves to avoid congestion in a self-organized, distributed manner. Still, diverse applications coexisting in the open Internet environment, lead to highly variable user demands and unpredictable resource availability.

Some developing distributed applications attempt to adapt themselves to Internet congestion dynamically in order to achieve desired temporal behavior in the unpredictable Internet environment. However, design and development of such network-adaptive distributed systems remain in an early stage. Research about the implication of Internet traffic characteristics for such distributed systems should prove timely and useful. Overlay networks, which are self-organized distributed systems created on top of the Internet, often need to monitor dynamic network conditions to achieve effective and stable application performance. Overlay networks have been designed to enable distributed applications to avoid congested paths, to support multimedia conferencing applications, and to manage caches for Web content delivery [1-3]. For example, Chu and colleagues [1] describe an application for multicast audio-video conferencing that can be transmitted across the Internet using a TCP (Transmission-Control Protocol) Friendly Rate Control (TFRC) protocol [4] in a self-organizing overlay network. The overlay network must adapt to Internet congestion by tracking available link bandwidth at an appropriate timescale. However, as Chu and colleagues state, there is an open issue associated with determining the most effective timescale for measurement and adaptation. Further studies, leading to a rigorous understanding of the characteristics of Internet congestion, may help resolve the issue.

Usually, successful adaptation based on monitoring at a specific timescale requires a strong correlation in measured conditions at that time granularity. Weaker correlation leads to larger errors in prediction, and thus less reliable feedback, since conditions fluctuate more rapidly than the time measurement interval. Some network protocols and algorithms use measurements of network state to guide future actions. For example, the transmission-control protocol (TCP) uses round-trip delay measurements to estimate when unacknowledged packets should be retransmitted, and selected admission control algorithms use measures of past load to predict future load [5]. These protocols and algorithms do not explicitly identify whether the relevant measures of network state vary sufficiently slowly to have a strong correlation over the timescales of interest for the intended controls.

Fortunately, network researchers have discovered evidence that time-series of measured Internet traffic usually exhibit long-range dependence (LRD) [6-8], defined as slowly decaying auto-correlation over a wide range of timescales. However, the LRD found in Internet traffic is believed to arise from highly variable user traffic, often represented with so-called heavy-tailed statistical distributions, and from TCP adaptation to network congestion [9]. For distributed applications that use TCP, the most prevalent transport protocol, adaptation to network congestion is built-in using specific algorithms; thus, network-adaptive applications have little to gain when operating over TCP. Similarly, since the Internet is a global shared resource, there is little that designers of network-adaptive applications can do to control the behavior of the many users sending traffic over the Internet for a wide variety of independent applications. Given these factors, designers of network-adaptive distributed systems usually choose to operate applications over a transport protocol other than TCP, and then to measure performance arising from the combined effects of TCP congestion adaptation (a factor in most current network applications) and global user behavior (which tends to generate heavy-tailed traffic patterns). Using such measurements, the end points in network-adaptive distributed systems adjust themselves to co-exist with their perceived operating environment. But over what timescales should network-adaptive distributed systems measure and adapt to global conditions?

In this paper, we analyze qualitatively the correlation structure of network traffic, which might have some relationship to the design of system-wide adaptation mechanisms intended to meet quality-of-service requirements in distributed systems. We believe that understanding the properties and causes of LRD in Internet traffic can help designers of network-adaptive distributed systems to produce more effective measurement and adaptation strategies. Where network traffic exhibits strong correlation over specific timescales, then monitoring and feedback control actions taken at those timescales should prove effective. Further, where strong correlation exists over larger timescales, network adaptation can occur more slowly, yet remain effective for a large distributed system. In such cases, it could also prove possible to achieve effective adaptation with lower measurement overhead.

We use simulation to illustrate effects on network congestion of some key factors: characteristics of offered traffic, end-to-end congestion-control mechanisms, link capacity, and network size. We apply wavelet-based analysis [10] to illustrate propagation of congestion over various timescales. We find that network size has a significant influence on the correlation structure of network traffic. Specifically, in a large network, correlation in congestion extends over a wide range of timescales, which implies that large network-adaptive, distributed applications executing over the Internet can rely on a robust and pronounced correlation structure to permit coarser-grained adjustments to global conditions. We also find that variations in load and available capacity, when occurring locally near application end points, can increase and decrease coupling with the global correlation structure. This suggests that network-adaptive distributed applications should vary the rate of adjustments rather than rely on a fixed adjustment rate.

The remainder of the paper is structured as three sections. Section II discusses our modeling and analysis approaches. We describe our representation of network structure, ON/OFF traffic sources, congestion-control algorithms, and routing. We also outline the wavelet-based analysis method. In Section III, we delineate our experiments and show related simulation results. Our experiments aim to distinguish among the effects of user behavior, transmission dynamics, and network structure. We present concluding remarks in Section IV.

II. Experimental Method

Simulating and understanding behaviors in a large-scale network present difficult challenges, arising from the need to search a large parameter space. However, if we restrict our simulation appropriately, then we can perhaps find a tractable approach that helps us to develop intuition and insight. We adopt a homogeneous topology that enables us to simulate a large-scale network, while varying selected factors: link capacity, traffic generation, transport mechanism, and network size. Using such a simulation, we can gauge change in a measure of interest when varying each factor, one by one.

A. *Modeling*. To provide a holistic view of network traffic, a model should encompass the variability and complexity inherent in a hierarchical network topology, including the effects of network-host interactions and the effects of protocol regimes and network controls. Though we restrict our attention to qualitative aspects of the correlation structure in network traffic, our model must still capture some important details of network structure, ON/OFF sources, congestion-control algorithms, and routing.

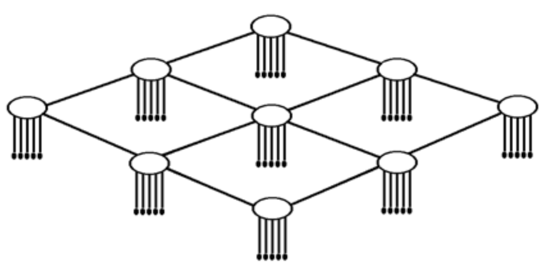

Figure 1: Model network structure

Network Structure. The network structure of our model is a two-tier topology: the upper tier for routers and the lower tier for hosts, as shown in Figure 1. In the upper tier, routers connect each other in a regular way to form a mesh-like (grid) network. An equal number of sources (n_s) attach to each router. Other hosts attached to the router take the role of receivers. When a source initiates a connection (ON period), a destination routing domain (i.e., router) is chosen randomly, and an available receiving host is assigned. The destination router must differ from the source router. A receiver is released when the source ends the connection. We limit the number of receivers for each routing domain to double the number of sources ($2n_s$). If a source selects a destination routing domain where all receivers are occupied, then another routing domain is selected randomly.

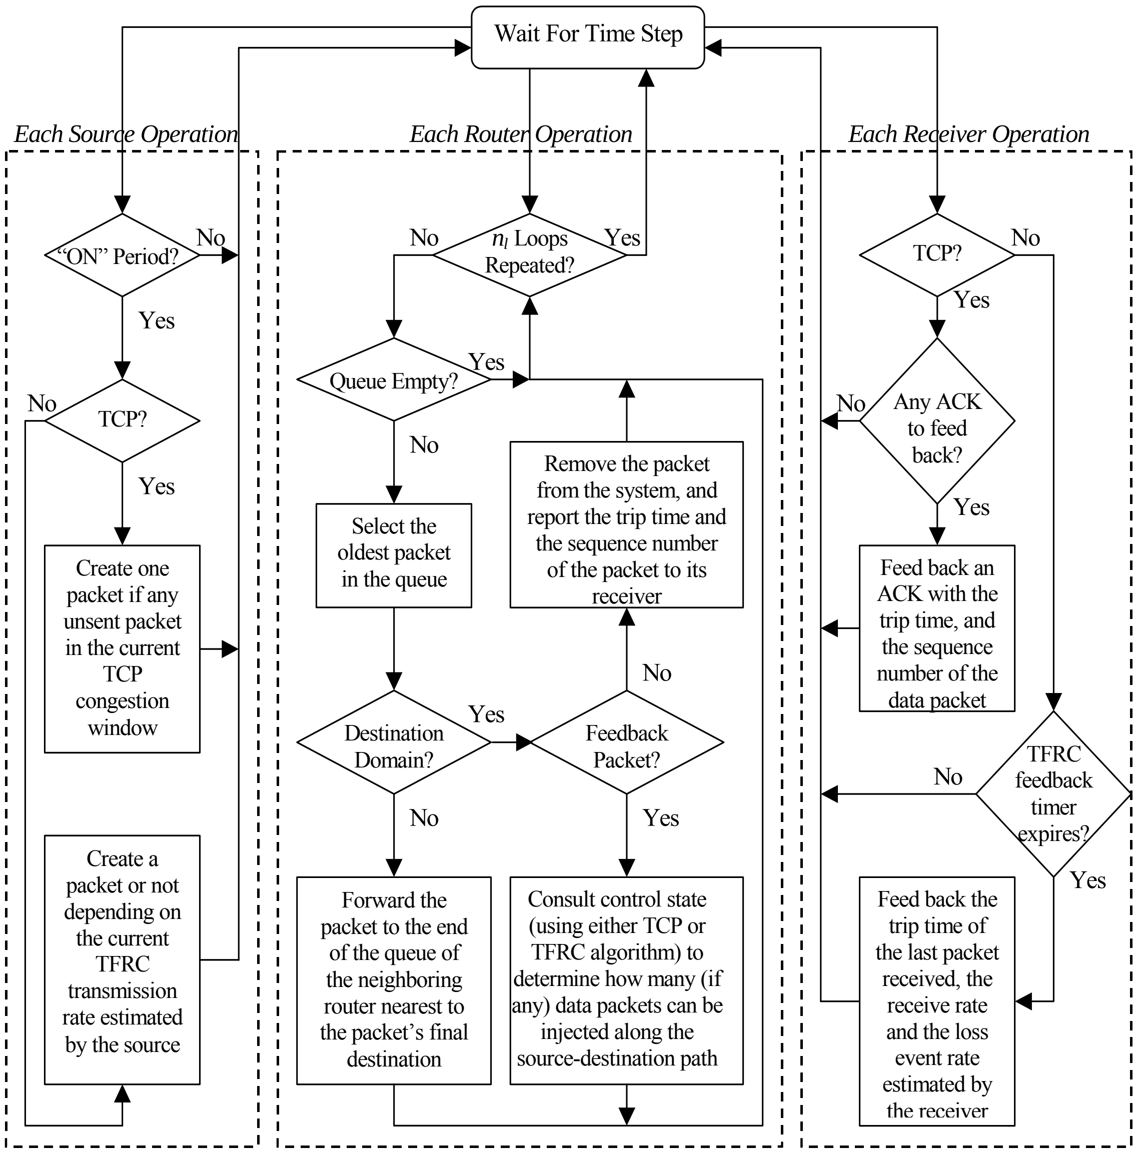

Figure 2: General parallel operations of sources, routers, and receivers at each time step

Our model operates at the packet level. To store and forward packets traveling between source-destination pairs, each router maintains a queue of limited length (50 packets), where arriving packets are stored until they can be processed: first-in, first-out. Packet movement occurs at discrete simulation time-steps. During each time-step, each source host can send at most one packet to its directly attached router. However, each router can forward multiple packets (n_l)

during each time-step. This forms a natural difference between two link types in the two tiers: the link capacity in the host tier is one packet-per-time-step (ppts), while the link capacity in the routing tier is n_l ppts. Figure 2 provides a schematic diagram of the parallel operations of sources, routers, and receivers at each time step.

We consider three parameters of network structure: number of sources (n_s), link capacity of the routing tier (n_l), and network size (L^2), where L is the number of routers along one side of the grid. In Figure 1, for example, $L = 3$ and $n_s = 5$; thus the network contains 9 routers and up to 45 simultaneously active connections. (Figure 1 omits receivers.)

Traffic Sources. Each source models traffic generation as an ON/OFF process, which alternates between wake and sleep periods. When awake, a source may send, subject to any restrictions imposed by congestion-control algorithms, one packet at each time-step to the source's first-hop router. The packet will be placed at the end of the router's queue. At the beginning of each "ON" period, a source randomly selects (uniform distribution) another routing domain (with an available destination host) as its sink. Each packet generated by the source during the same "ON" period has the same destination address. When sleeping, the source does not generate new packets at each time-step. ON/OFF sources provide a convenient model of user behavior.

We modulate the frequency and duration of ON/OFF periods by selecting arrival and departure times according to various statistical distributions. Here, we use one of two distributions. Most empirical measurements on the Internet observe a heavy-tailed distribution of transferred file sizes [8]. Some researchers believe long-range dependence arises in the Internet from the high user variability represented by such heavy-tailed distributions [11]. To investigate this belief, we need a distribution to represent such user behavior, and another distribution to represent lower user variability. In this latter case, we represent wake and sleep period durations as an exponential process with parameters λ_{on} and λ_{off}.

To model high user variability, we represent the wake and sleep period durations using the Pareto distribution, which is frequently used to model the heavy-tailed characteristic of Internet file transfers. The Pareto distribution function has the following form: $P\,[X \le x] = 1 - (k/x)^{\alpha}$, $k \le x$, where $0 < \alpha < 2$ is the shape parameter. Here, we use $\alpha = 1.2$. The mean is given by $\alpha k/(\alpha - 1)$. In this paper, we sometimes mix the Pareto and exponential distributions, using a Pareto "ON" period, and an exponential "OFF" with λ_{off}. To keep the same average "ON" duration for both distributions, we let $k = (\alpha - 1)\lambda_{on}/\alpha$.

Congestion Control Algorithms. To achieve traffic dynamics across all timescales, our model implements a parallel system at the packet level, where packets transit along connections between source-destination pairs. Each connection operates full duplex under one of two traffic-control regimes: TCP and TCP friendly rate control (TFRC) [12].

Modern TCP implementations contain four intertwined algorithms: slow start, congestion avoidance, fast retransmit, and fast recovery. Variants of TCP include Tahoe, Reno, NewReno, and SACK TCP; the last three differ only in their response when multiple packets are dropped from a widow of data [13]. In this paper, we model Reno TCP, except that our model reduces the congestion window to half the current window size after receiving one, instead of three, duplicate ACKs for the same packet. While packets can be lost in our model, all packets on a connection take the same route, so no packet reordering occurs.

TFRC, relative to TCP, has a more smoothly varying transmission rate. The corresponding cost is a slower response to transient changes in congestion or to sudden increases in available bandwidth [14]. TFRC uses a receiver-driven mechanism, where the receiver calculates congestion-control information—i.e., the loss rate—and feeds back to the sender. The sender also uses these feedback messages to measure the round-trip time (RTT). The sender inputs the loss rate and RTT into a throughput equation, which yields an acceptable transmission rate. Then, the sender in our model adjusts its interval between packet transmissions to match the acceptable rate.

Whether used in TCP or TFRC, each packet in our model carries several pieces of information: source address (router and host), destination address (router and host), creation time, and sequence number. In addition, the sender and receiver on each connection maintain state information, and exchange information via packets. In particular, for TCP the receiver maintains the expected sequence number to identify if a packet is lost, and to inform the sender when a packet drop occurs. For TFRC, the receiver maintains estimates of the loss rate and the packets-received rate, and records the timestamp of the last packet received. The receiver periodically sends this information, along with an estimated RTT, to the sender.

Routing. Each source-destination pair has a constant, shortest path. Instead of maintaining a forwarding table for each router, we compute a routing for each packet. To select the proper next-hop for a packet, the forwarding router computes the distance from each of its four neighboring routers to the packet's destination router. Given the regular grid topology of our model, distance calculations can be performed easily. Following Fuks [15], we use a metric defined for models with a periodic boundary condition to determine the distance between routers. Where multiple neighboring routers prove equidistant from the destination (at most two choices in our model), we consistently choose the left direction, which provides a constant path for all packets on a connection. Our routing technique leads to a constant, shortest path for each source-destination pair in one direction, and a complementary shortest path for the reverse direction.

B. *Analysis Technique*. Characterizing network traffic requires knowledge about the statistical properties of packet arrivals over various timescales. Empirical network traffic exhibits LRD at large timescales, as manifested by slowly decaying autocorrelations. LRD is one of several equivalent ways to describe second-order self-similarity [11], but we focus on a more general concept: correlation structure of network traffic.

Correlation structure can be observed through the autocorrelation function, or its Fourier transform—i.e., the power spectral density. As indicated by Figueiredo and colleagues [16], the periodogram method exhibits possible pitfalls. On the other hand, a wavelet-based technique [10], which is frequently used to analyze long-range dependent data and to estimate the associated Hurst parameter, provides a natural and effective tool to reveal the correlation structure of network traffic across a wide range of timescales. Because the wavelet transform divides data into different frequency components and analyzes each component with a resolution matched to its scale, the coefficients of wavelet decomposition can be used directly to study the scale (or frequency) dependent properties of data. The discrete wavelet transform represents a time series X of size N at a scaling level j by a set of wavelet coefficients $d_X(j, k)$, $k = 1, 2, \ldots, N_j$, where $N_j = 2^{-j}N$. The coefficient $|d_X(j, k)|^2$ measures the amount of energy in X about the time $2^j k$ and about the frequency $2^{-j}f_0$, where f_0 is determined by the sampling rate of the time series and the choice of the analyzing wavelet. Average energy at scale j (where scale j is referred to as an octave) is the average of the sum of the squared wavelet coefficients $|d_X(j, k)|^2$; i.e., $E_j = \frac{1}{N_j}\sum_k |d_X(j,k)|^2$.

As indicated by Figueiredo and colleagues [16], E_j is really an estimate of the power spectral density about the frequency $2^{-j}f_0$. We can plot $\log_2 E_j$ as a function of scale j, from finest to coarsest scales, and investigate the structure in the energy scale plot[1].

III. Experiments and Simulation Results

Other researchers [8, 9, 16-18] have investigated some aspects of user and network behavior that contribute to different characteristics in the dynamics of network traffic. We attempt to verify some of the existing results, and we study effects of traffic sources, transport mechanisms, and network structure.

[1] Matlab code for the wavelet method is available from http://www.emulab.ee.mu.oz.au/~darryl/.

A. *Effects of the Application Layer*. Using an abstract ON/OFF model to mimic user behavior (a property of the application level), we consider the effect of three parameters: mean values of ON/OFF durations (λ_{on} and λ_{off}), and, when using the Pareto distribution for the "ON" period, the shape parameter α. We first investigate the following network configuration: network size L = 3, number of sources n_s = 10, and link capacity n_l = 5. We use TCP as the transport level. The application level comprises exponentially distributed ON/OFF periods with λ_{on} = 200 and λ_{off} = 2000.

Starting with a random initial condition, after discarding a transient period of 10^4 time-steps, we analyze the traffic on one link between routers in one direction. We plot the resulting correlation structure as $y_j = \log_2 E_j$ vs. j (the top plot in Figure 3). In this paper, we record enough data to yield an energy scale plot that spans 20 octaves. Note that the finest scale description of traffic dynamics depends on the selection of n_l. We focus solely on the large timescale features, checking for a more or less linear relationship.

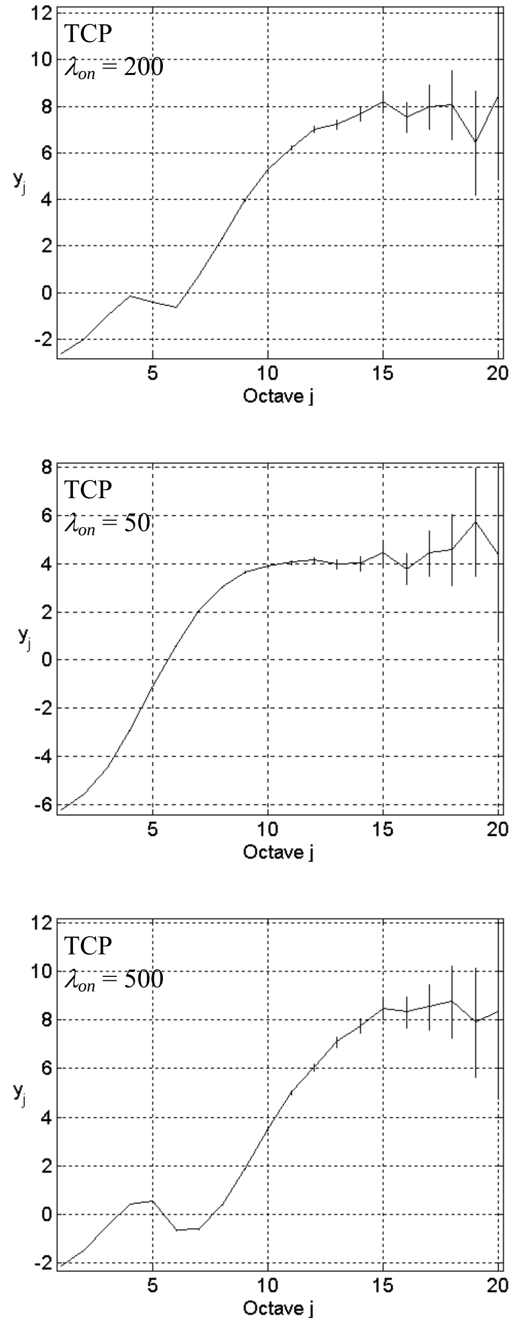

Figure 3: Energy scale plots for L = 3, n_s = 10, n_l = 5, TCP in transport level, the exponential distribution ON/OFF with λ_{on} = 200 (top), 50 (middle), 500 (bottom), and λ_{off} = 2000

In the top plot of Figure 3, we observe a strong correlation structure that spans around 5 octaves. The curve departs from linearity at the mediumto-small time scale ($j < 6$), and becomes flat after $j \geq 11$. The linear part of the curve implies that the autocorrelation function decays according to a power law within a limited range of time scales. The flat part of the curve indicates that the autocorrelation function decays exponentially. Obviously the correlation structure is not consistent with LRD. Note, however, that the correlation structure is consistent with similar results reported by Figueiredo and colleagues [16, 17].

To investigate interaction between network congestion and offered traffic, absent high user variability, we keep all parameters fixed except for λ_{on}, where we use 50 and 500, shown in the middle and bottom, respectively, energy-scale plots of Figure 3. We can see that changing λ_{on} alters the correlation structure. Specifically, as λ_{on} increases, the low-frequency energy (right portion of the energy-scale plot) tends to increase, and the linear part extends over a wider range of time scales. A similar effect also appears when changing λ_{off}.

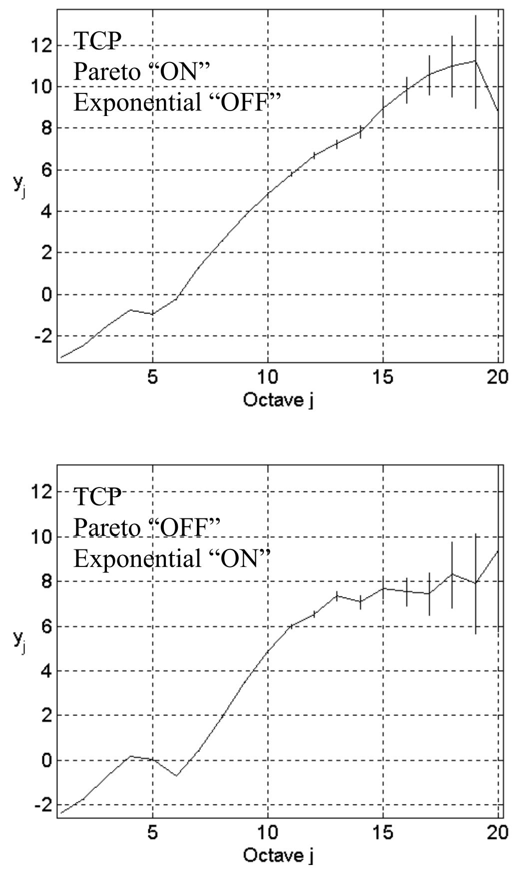

Figure 4: Energy scale plots for L = 3, n_s = 10, n_l = 5, λ_{on} = 200, and λ_{off} = 2000, TCP in transport level, and the Pareto distribution "ON" (α = 1.2) with the exponential "OFF" (top), or the Pareto distribution "OFF" (α = 1.2) with the exponential "ON" (bottom)

To consider the effect of high user variability, we keep the same network configuration except for the "ON" period, where we use the Pareto distribution (λ_{on} = 200, α = 1.2, and λ_{off} = 2000). The top plot in Figure 4 gives the corresponding energy-scale plot, which verifies previous results [8, 9, 17, 18] that heavy-tailed file sizes lead to LRD over large time scales. Varying λ_{on} and λ_{off}, we find this characteristic is robust. However, we cannot find this characteristic if we use a Pareto distribution to model heavy-tailed idle time. For example, keeping the same parameters, except using exponential "ON" and Pareto "OFF" (λ_{on} = 200, and λ_{off} = 2000, α = 1.2,), the energy-scale plot in the bottom of Figure 4 exhibits a flattening structure at low frequencies.

In summary, high variability in file sizes can result in a strong correlation structure over a wide range of timescales, while low variability yields correlation over a more limited range. The correlation structures arise from many connections interacting under dynamic conditions, which can be shaped by offered load. Others [16] have suggested that two mechanisms inside TCP—exponential timeout back-off and congestion avoidance—contribute to the correlation structure. Will the same correlation structure arise when transporting data using a congestion-control mechanism without one or both of these TCP mechanisms? We investigate this question next.

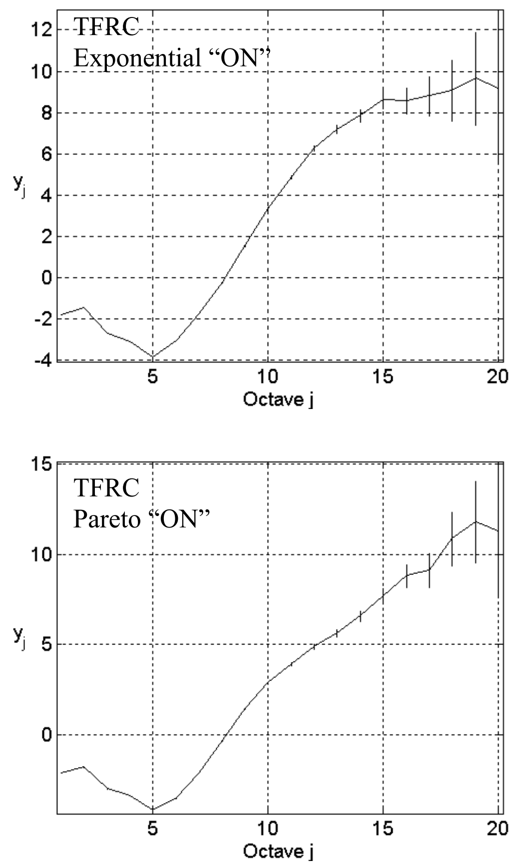

Figure 5: Energy scale plots for L = 3, n_s = 10, n_l = 5, λ_{on} = 200, and λ_{off} = 2000, TFRC in transport level, and the exponential distribution ON/OFF (top), or the Pareto distribution "ON" (α = 1.2) with the exponential "OFF" (bottom)

B. *Effects of the Transport Layer*. We set up the network configuration with the same parameters used to obtain the top plot in Figure 3, except that TFRC replaces TCP. The energy scale plot appears as the top graph in Figure 5. We find that the correlation structure at small time scales differs from the case of TCP (top plot in Figure 3). However, similar to TCP, a strong correlation structure spans a limited range of octaves (j = 6 to 12). We also observed (not shown here) that offered traffic has similar effects on the correlation structure, whether using TCP or TFRC. So it seems that the limited strong correlation structure does not rely on particular transport mechanisms. Congestion feedback algorithms, other than those used in TCP, may also yield a limited strong correlation structure as a result of interactions among many connections.

What about the effects of high user variability? Will heavy-tailed file sizes lead to LRD under TFRC? Other researchers [9, 19] believe that, under TCP or flow-controlled UDP, LRD of aggregate traffic will appear as long as connection durations or object sizes being transported are heavy-tailed. In other words, they believe that LRD in aggregate traffic is insensitive to details in the protocol stack or the network configuration. Despite such beliefs, we have not found any research that investigates this question using any transport protocol other than TCP or an open-loop flow-controlled unreliable transport protocol, used by Park and colleagues [19]. We investigate the question with TFRC.

Keeping the same network configuration (see Figure 5), we substitute the Pareto distribution "ON" (λ_{on} = 200, α = 1.2, and λ_{off} = 2000) in place of exponential "ON". From the correlation structure in the corresponding energy-scale plot, which appears as the bottom graph in Figure 5, heavy-tailed file sizes appear to give rise to LRD over large time scales under TFRC.

In the next experiment, we change only one parameter, reducing the link capacity to n_l = 2. The corresponding energy scale plot (top graph in Figure 6) does not exhibit the expected correlation structure. For TFRC, shrinking link capacity destroys the LRD structure induced by heavy-tailed file-size distributions. In contrast, when substituting TCP for TFRC (bottom plot in Figure 6), retaining all other parameter settings including the reduced link capacity, we find, as indicated elsewhere [9], that TCP maintains the LRD structure introduced by heavy-tailed file sizes. This suggests that LRD of aggregate traffic might *not* be insensitive to details in the protocol stack or network configuration, which motivates us to explore the effects of network structure.

C. *Effect of Network Structure: Relative Bandwidth*. We study the effect of the network structure on traffic dynamics by modulating three parameters: number of sources (n_s), link capacity (n_l), and network size (L). We first identify how shrinking or expanding link capacity influences the correlation structure. We set L = 3, n_s = 10, and n_l = 2 or 20; we use TCP or TFRC as the transport level; we use exponential ON/OFF (λ_{on} = 200 and λ_{off} = 2000) in the application level. Figure 7 provides the energy-scale plots from the simulation results.

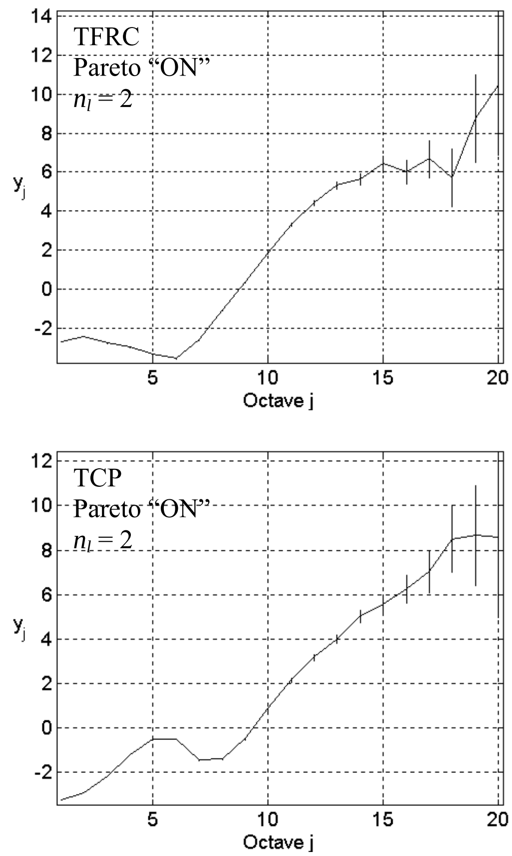

Figure 6: Energy scale plots for L = 3, n_s = 10, n_l = 2, λ_{on} = 200, and λ_{off} = 2000, the Pareto distribution "ON" (α = 1.2) with the exponential "OFF", and TFRC (top), or TCP (bottom) in transport level

The top row of plots in Figure 7 depict correlation structures when n_l = 2 with either TCP (left) or TFRC (right) used as the transport layer. The bottom row of plots shows correlation structures arising when n_l = 20 with a transport layer of either TCP (left) or TFRC (right). Recall that, since the basic simulation time step is constant, the finest description of traffic dynamics relies on the selection of n_l. The effects of this fact can be observed in Figure 7, where the bottom row of plots gives a coarser description for the finest time scale than do the top row of plots.

Examining the correlation structure at large time scales, we find that changing link capacity (n_l) alters the correlation structure. Reducing link capacity tends to strengthen the correlation structure, while expanding link capacity loosens the degree of dependence in the traffic. A similar effect appears through changing n_s.

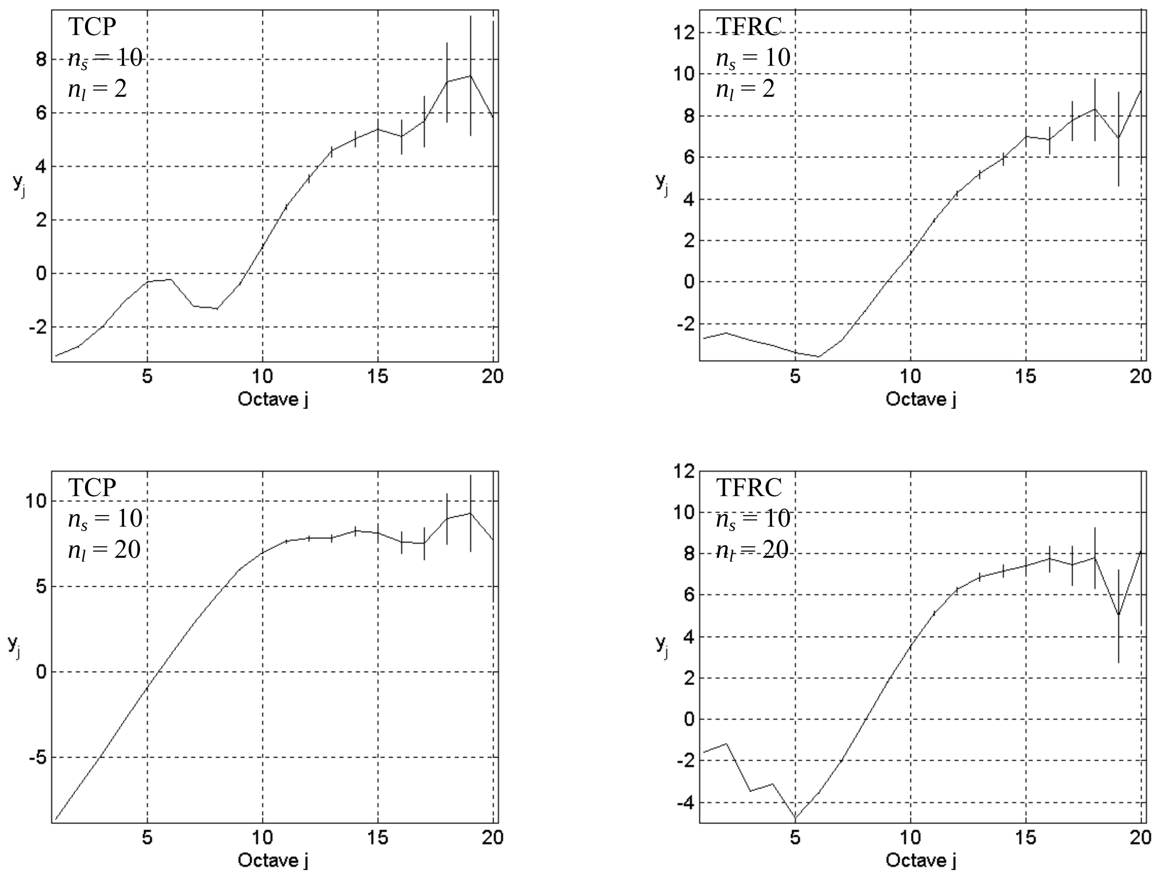

Figure 7: Energy scale plots for $L = 3$, $n_s = 10$, $n_l = 2$ (top), or 20 (bottom), the exponential distribution ON/OFF with $\lambda_{on} = 200$ and $\lambda_{off} = 2000$, and TCP (left) or TRFC (right) in transport level.

Retaining the same network configuration, we set $n_s = 40$, and $n_l = 5$. In Figure 8, we show the corresponding energy-scale plots for TCP (top left) and TFRC (top right). Comparing against the top plot in Figure 3 (TCP with $n_l = 5$) and the top plot in Figure 5 (TFRC with $n_l = 5$), we find that increasing n_s tends to strengthen the correlation structure. The effect is similar to effects from regulating link capacity and varying $\lambda_{on}/\lambda_{off}$. To understand this relationship, we also show, in the bottom of Figure 8, two additional energy-scale plots: one for TCP (bottom left) and one for TFRC (bottom right). The parameters for these plots correspond to the same parameters used in the top row of plots in Figure 8, except that we increase the link capacity from $n_l = 5$ to $n_l = 20$. With the increased link capacity, the correlation structures in each plot seem to return to their

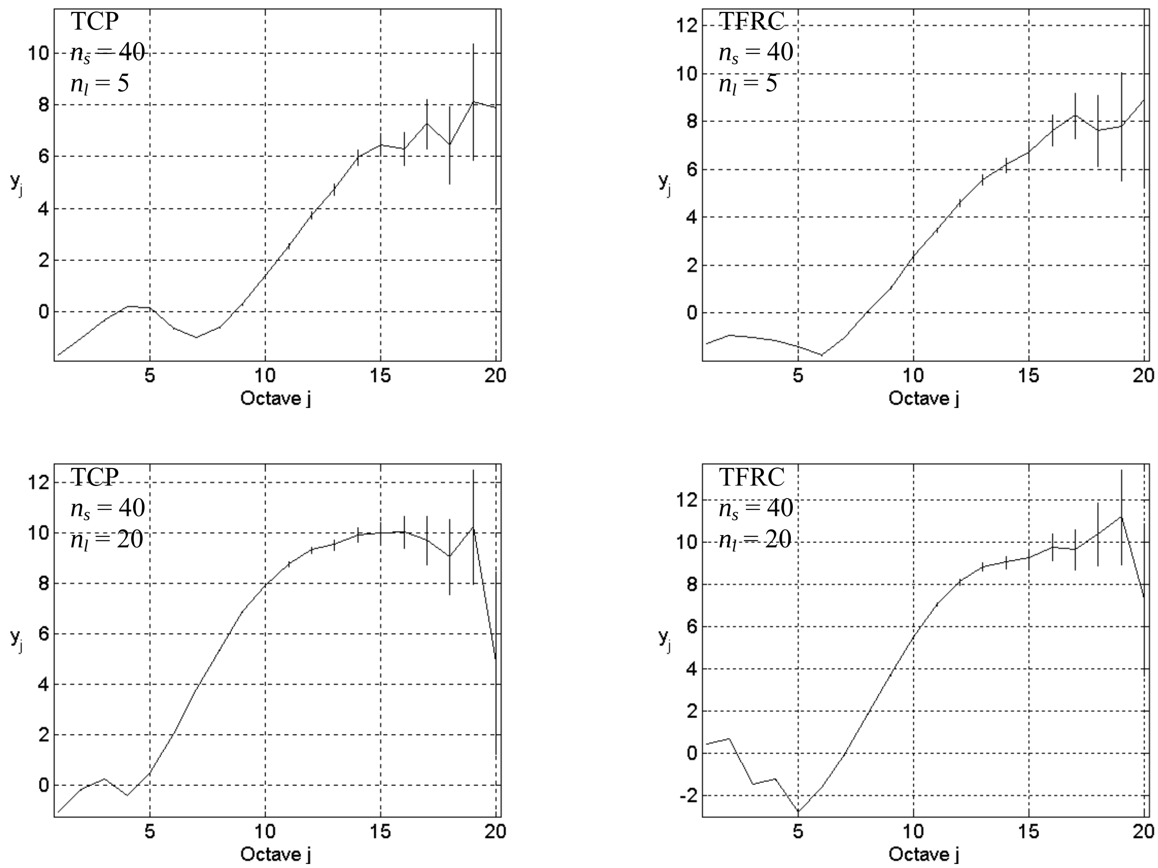

Figure 8: Energy scale plots for $L = 3$, $n_s = 40$, $n_l = 5$ (top), or 20 (bottom), the exponential distribution ON/OFF with $\lambda_{on} = 200$ and $\lambda_{off} = 2000$, and the TCP (left) or the TRFC (right) in transport level

original shapes, as depicted in the top plots in Figure 3 (for TCP) and Figure 5 (for TFRC).

Offered traffic (represented by λ_{on}, λ_{off}, and n_s) and shared network capacity combine to act as traffic-shapers, strengthening and loosening correlation structure, which can be offset by congestion-control mechanisms. This view might help to explain why TCP, and its variants, are prone to instabilities when combined with increases in network capacity [20]. For network-adaptive distributed systems, an increased correlation structure might provide more stable traffic patterns that permit network measurements, and subsequent behavioral adjustments, to occur less frequently. However, in case of correlation degradation, network-adaptive systems may have to vary measurement pace to match the more rapid fluctuation.

On the other hand, more surprising interactions might exist in a large-scale network, where the utilization of network capacity could be more influenced by spatial relationships. We investigate this next.

D. *Effect of Network Structure: Network Size*. A key property of the Internet is its large scale. In July 1998, as reported by Cowie and colleagues [21], the Internet comprised a collection of about 4,000 interconnected routing domains (or autonomous systems). Does network size play a significant role in defining the correlation structure of network traffic?

Even if our model has the potential to answer this question, absent a high-speed parallel computer system, we may still spend a very long time to simulate a network with larger size, capacity, and number of hosts, and to collect sufficient data for wavelet-based analysis. To surmount this obstacle, we chose to investigate a situation where network sources can congest the network backbone, though this is counter to the conventional case where network congestion appears more frequently on network access links. Since our investigation considers comparative rather than absolute results, our limiting assumptions might lead to an informative outcome while reducing computational requirements. We set up our experiments using the following parameters: $n_s = 5$, $n_l = 1$, and $L = 3$, 9, and 27; TCP or TFRC in transport level; exponential distribution ON/OFF with $\lambda_{on} = 200$ and $\lambda_{off} = 2000$ in application level. The energy-scale plots appear in Figure 9 for TCP (left column) and TFRC (right column), with network size increasing from top ($L = 3$) to bottom ($L = 27$). These plots show that as the network size increases, the flat portion of the curve (indicating exponential decay in the autocorrelation function) diminishes little by little, while the linear portion seems gradually to increase in significance. So it might be reasonable to expect the linear part to extend over some large time scales as the network reaches Internet size, even without high variability in user behavior. This suggests that network size has greater influence than other factors on network congestion.

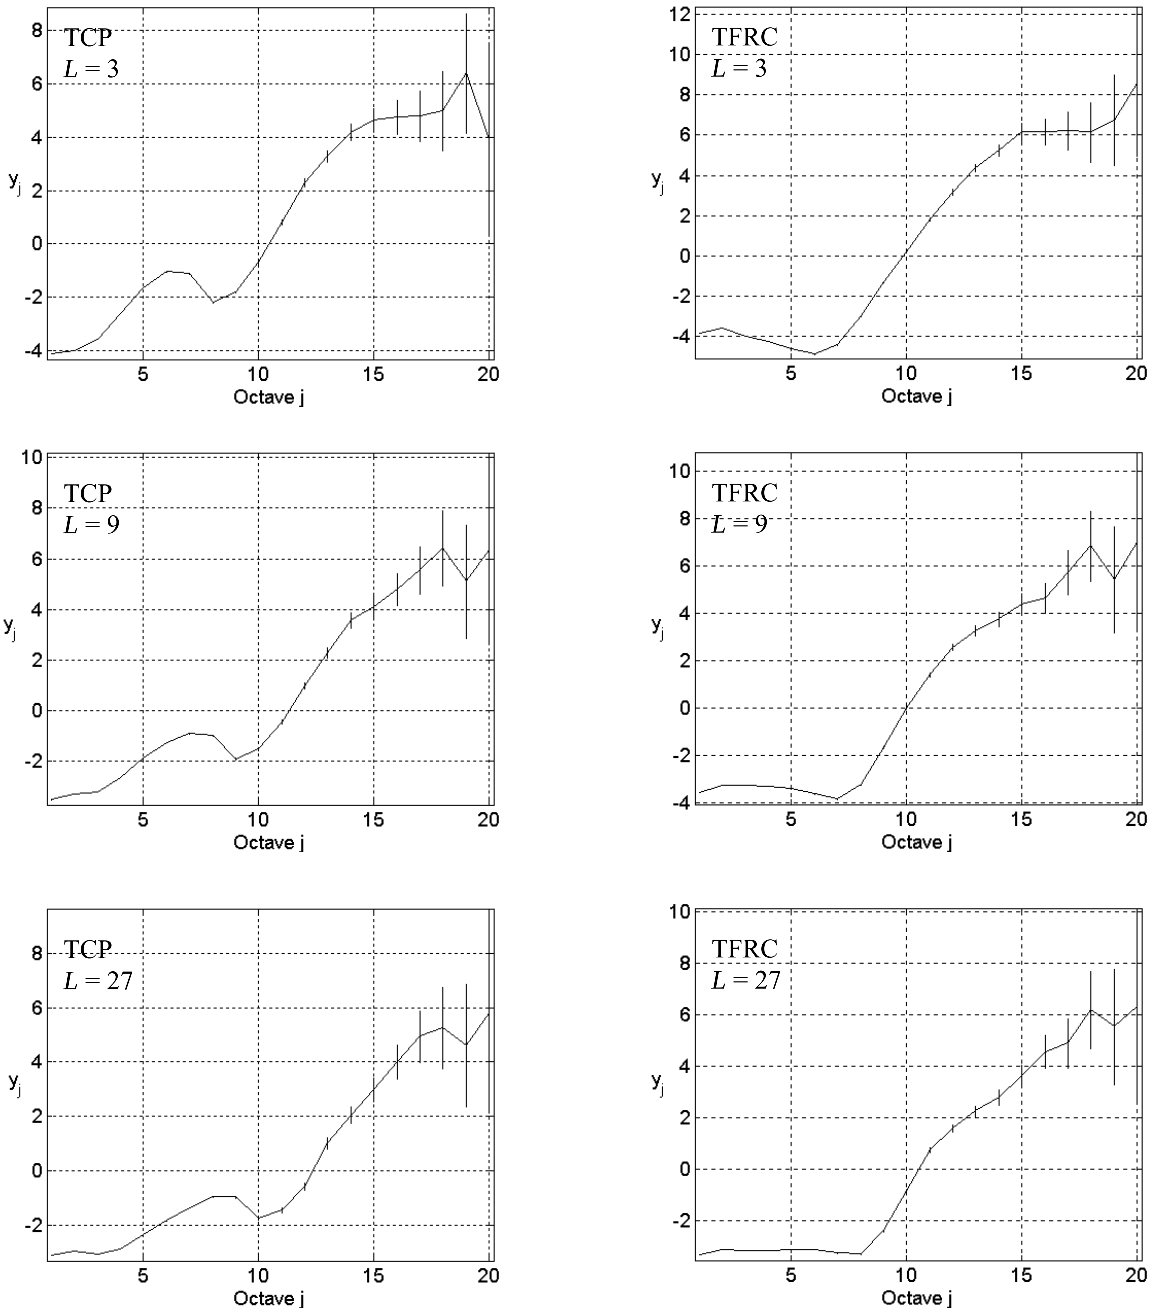

Figure 9: Energy scale plots for $L = 3$ (top), 9 (middle), and 27 (bottom), $n_s = 5$, $n_l = 1$, the exponential distribution ON/OFF with $\lambda_{on} = 200$ and $\lambda_{off} = 2000$, and TCP (left) or TRFC (right) in transport level

IV. Conclusions

By focusing our simulation model on describing comparative rather than absolute behavior, we conducted a systematic search, using wavelet-based analysis, to identify and understand significant phenomena influencing the correlation structure of network congestion. Our findings imply that observed traffic characteristics derive from combined effects, which might be overlooked as researchers search for invariants from empirical data. We also reported several new observations and recommendations with significant implications for traffic characterization, and for the design of network-adaptive distributed applications.

First, we find that congestion-control mechanisms affect the characteristics of timescale dynamics in network traffic. While both TCP and TFRC can produce strong correlation over a limited range of timescales, their influence on the correlation structure might differ, especially at large timescales, where TFRC sometimes fails to maintain the LRD structure induced by high user variability. In effect, congestion-control mechanisms might balance offered load and network capacity, leading to invariants over a limited range of timescales. These invariants appear not to hold across all timescales.

Second, we find that the correlation structure of traffic should be controllable by modulating available bandwidth. Since a network is a driven system, the correlation structure of network traffic arises from significant interaction between offered load and network capacity. Thus changes in offered traffic change the correlation structure. From the perspective of Internet traffic engineering, we can imagine that changing available capacity to follow the daily pattern of traffic could shape the correlation structure. However, from the view of network-adaptive distributed systems, we see the need to vary the pace of adaptation to account for more rapid fluctuation as correlation structure degrades. An increased correlation structure might provide distributed systems with more stable traffic patterns that permit network measurements, and subsequent behavioral adjustments, to occur less frequently.

Third, we find that a similar correlation structure to that seen for measured Internet traffic may arise in very large networks, even without high user variability. In large networks, much of the traffic must travel through multiple domains en route from source to destination. The observed correlation structure of traffic arises from the collective effect of all transit flows. Where concurrent connections share one link in a larger network, the longer-distance connections, which need more time for feedback control, must be responsible for the larger timescale correlation structure because connections using either TCP or TFRC can themselves exhibit strong correlation over a limited range of timescales. Therefore, the spatial span of connections appears to be a significant influence on correlation structure. Large distributed applications might be able to exploit this general characteristic to achieve network-wide adaptation. To support this insight, we conducted a further study that demonstrates an ability to monitor shifting network-wide patterns at some timescales [22].

Finally, effective design of network-adaptive distributed systems should benefit from an understanding of network-wide traffic flow. While researchers have provided an understanding of the properties of various quality-of-service mechanisms on a local scale, the effects of such mechanisms have not been studied on a global scale in large networks, where cause-effect relationships might not be inferred readily from the behavior of individual elements. In-depth understanding of the spatiotemporal variance of network traffic should be viewed as an important research direction that could lead eventually to a solid methodology for designing network-adaptive distributed systems.

V. References

[1] Y. Chu, S. Rao, S. Seshan, and H. Zhang, Enabling Conferencing Applications on the Internet using an Overlay Multicast Architecture, In: Proceedings of ACM SIGCOMM'01, San Diego, CA, USA, Aug. 2001.

[2] D. G. Andersen, H. Balakrishnan, M. F. Kaashoek, and R. Morris, Resilient overlay networks, In: Proceedings of SOSP '01, Banff, Canada, Oct. 2001.

[3] J. Jannotti, D. Gifford, K. Johnson, M. Kaashoek, and J. J.W. O'Toole, Overcast: Reliable Multicasting with an Overlay Network, In: 4th Symposium on Operating Systems Design and Implementation (OSDI), San Diego, CA, USA, Oct. 2000.
[4] S. Floyd, M. Handley, J. Padhye, and J. Widmer, Equation-Based Congestion Control for Unicast Applications, in: Proc. SIGCOMM 2000, August 2000.
[5] L. Breslau, S. Jamin, and S. Shenker, Comments on the performance of measurement-based admission control algorithms, In: Proc. IEEE INFOCOM 2000, Tel Aviv, Israel, March 2000.
[6] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, On the self-similar nature of Ethernet traffic, in: Proc. ACM SIGCOMM ’93, 183-193, 1993.
[7] V. Paxson and S. Floyd, Wide-area traffic: The failure of Poisson modeling, in: Proc. ACM SIGCOMM ’94, 257-268, 1994.
[8] M. E. Crovella and A. Bestavros, Self-similarity in world wide web traffic: Evidence and possible causes, in: Proceedings of the 1996 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems, May 1996.
[9] K. Park, G. Kim, and M. Crovella, On the relationship between file sizes, transport protocols, and self-similar network traffic, in: Proc. International Conference on Network Protocols, 171-180, October, 1996.
[10] P. Abry and D. Veitch, Wavelet analysis of long-range dependent traffic, IEEE Transactions on Information Theory, 44, pp. 2-15, 1998.
[11] W. Willinger, M. S. Taqqu, R. Sherman, and D. V. Wilson, Self-similarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level, in: Proc. ACM SIGCOMM ’95, 100-113, 1995.
[12] M. Handley, J. Pahdye, S. Floyd, and J. Widmer, TCP Friendly Rate Control (TFRC): Protocol Specification, draft-ietf-tsvwg-tfrc-03.txt, Internet draft, work-in-progress, May 2001.
[13] K. Fall, S. Floyd, Simulation-based comparisons of Tahoe, Reno, and SACK TCP, Computer Communication Review, 26 (3), pp. 5-21, 1996.
[14] D. Bansal, H. Balakrishnan, S. Floyd, S. Shenker, Dynamic behavior of slowly-responsive congestion control algorithms, in: Proc. ACM SIGCOMM, San Diego, CA, September 2001.
[15] H. Fuks, A. T. Lawniczak, Performance of data networks with random links, Mathematics and Computers in Simulation, Vol. 51, pp.101-117, 1999.
[16] D. R. Figueiredo, B. Liu, V. Misra, and D. Towsley, On the Autocorrelation Structure of TCP Traffic, Computer Networks Journal Special Issue on Advances in Modeling and Engineering of Long-Range Dependent Traffic, 2002.
[17] A. Feldmann, A. C. Gilbert, P. Huang, and W. Willinger, Dynamics of IP traffic: A study of the role of variability and the impact of control, in: Proc. ACM SIGCOMM ’99, 301-313, 1999.
[18] A. Feldmann, A. C. Gilbert, W. Willinger and T. G. Kurtz, The changing nature of network traffic: Scaling phenomena, ACM SIGCOMM Computer Communication Review, Volume 28, No. 2, April 1998, pp. 5-29.
[19] K. Park, G. Kim, and M. Crovella, On the effect of traffic self-similarity on network performance, In: Proc. SPIE International Conference on Performance and Control of Network Systems, pp. 296-310, 1997.
[20] F. Paganini, J. C. Doyle and S. H. Low, Scalable laws for stable network congestion control, in: Proceedings of the IEEE Conference on Decision & Control, pp.185-190, Orlando, FL, December 4-7, 2001.
[21] J. H. Cowie, D. M. Nicol and A. T. Ogielski, Modeling the global Internet, Computing in Science & Engineering, Vol. 1, No. 1, pp. 42-50, January/February 1999.

[22] J. Yuan, K. Mills, A cross-correlation based method for spatial-temporal traffic analysis, submitted, 2003.